

NETWORK ADDRESS TRANSLATION GATEWAY FOR LOCAL AREA NETWORKS USING LOCAL IP ADDRESSES AND NON-TRANSLATABLE PORT ADDRESSES

BACKGROUND

Virtual private networking (VPN) using TCP/IP enables secure, high speed communications between remote computing sites, using the internet as a communications medium. Information passing between sites across the internet may be protected against interception by unwanted eavesdroppers or malicious hackers by a variety of security measures. Effective security measures must, at a minimum, incorporate functions that ensure any or all of the following protections: Data integrity against the inadvertent or malicious modification of data during transit; prevention against denial-of-service attacks by employing anti-repeat measures; source authentication; confidentiality of source address and other header information during transit; and packet payload protection against unwanted interception. One standard model for providing internet security is the Internet Protocol Security suite, IPSec. IPSec works with the TCP/IP communications protocol to provide secure communications between devices connected to the internet or connected to private LANs (Local Area Networks) that are, themselves, connected to the internet.

The TCP/IP (Transmission Control Protocol/Internet Protocol) protocol suite uses IP addresses to identify each device on a network. A global IP address uniquely identifies a device on the internet. Such devices may be computers, printers, routers, switches, gateways, or other network devices. Devices having global IP addresses may be directly referenced as a source or destination on the internet. The TCP/IP communications protocol, however, is not exclusively limited to the internet, but may be used on private LANs as well. Private LANs using TCP/IP frequently use "local" IP addresses for network devices. Although no two devices on a private LAN may share the same local IP address,

private LANs are isolated from the internet, and local devices on the LAN are not visible from the internet. IP addresses for local devices therefore need not be “globally” unique. A LAN using local IP addresses will be connected to the internet through a gateway, which is a device that may filter or route messages between the LAN and the internet. Since the gateway is directly attached to the internet, and is visible to it, the gateway must have a globally unique IP address for communicating across the internet. However, since the LAN is not directly visible from the internet, local machines on the LAN do not require IP addresses that are globally unique.

TCP/IP is the communications protocol that is used on the internet. Information to be communicated using TCP/IP is contained in “datagrams.” A datagram consists of a discrete “packet” of information to which one or more headers are appended. Headers contain information needed by TCP/IP to direct the packet to its intended destination and to ensure its proper handling during transit. Each datagram is individually addressable, and may be a connection-oriented TCP datagram or a “connectionless” UDP (User Datagram Protocol) datagram. Each UDP datagram includes an IP header and a UDP header. The IP header contains at least a “source” IP address and a “destination” IP address, while the UDP header includes source and destination service addresses (port addresses, given as numbers). In IPv4, IP addresses are 32 bits in length, and are associated with the now-familiar xxx.xxx.xxx.xxx format. In this format, each three-digit segment is a binary octet representing a number between 0 and 255. A complete IP address combines the address of a logical network or network segment with the address of a “node” (device) on the network. The address of the network or network segment may encompass the first 3, 6, or 9 digits of the IP address. A device on the network or network

1 segment is identified by a node address that consists of those remaining digits that are not
2 used in the network or network segment address.

3 The source and destination service addresses contained in a UDP header are 16-bit
4 numbers, variously known as "ports" or "sockets," which are used to direct the packet to
5 an intended process that is active on the sending or receiving device. The term "port," or
6 "port address," as used herein, refers to a service address field in a UDP header. Although
7 in theory there are as many ports as there are addresses in a 16-bit number, by convention
8 many port addresses have been reserved for established processes. Thus, for example,
9 port 80 is reserved for HTTP, and ports 20 and 21 are reserved for FTP. Through the use
10 of port addresses, data arriving at a local machine running more than one process will be
11 directed to the process for which it was intended. Where a process running on a local host
12 is not one of the reserved processes, the local host may select any port number from a
13 pool of unreserved port numbers to identify the "source" process. A reply packet
14 referencing that port number in the "destination" field will be directed to the process.

15 With the explosive growth of internet usage during the last decade, and its projected
16 growth in the future, globally unique IP addresses have become a scarce resource. In
17 addition, many businesses maintaining private LANs have little or no need for each
18 computer and device on the LAN to have a unique global IP address. Many such
19 businesses would, in any event, prefer to maintain the confidentiality of their computers'
20 IP addresses. Rather than waste limited global resources by giving each local device a
21 unique global IP address, many private LANs utilize local IP addresses for devices on the
22 LAN. In order to provide connectivity to the internet, such LANs will employ a single

globally unique address to be used on the internet by the gateway separating the LAN from the internet.

Through the use of Network Address Translation (NAT) techniques, a gateway device separating a LAN from the internet can provide security as a firewall while enabling machines with local IP addresses to access the internet through the unique global address of the gateway. A device on a LAN may have a static local IP address, or it may have a local IP address dynamically assigned to it at log on. The gateway maintains a translation table with the local IP addresses for each device on the LAN. A UDP packet sent from a local machine and destined for the internet will have its local IP address and port address identified in the source fields of the IP and UDP headers, respectively. The gateway will receive the packet from the local machine, will substitute its external globally-unique IP address and a new port address (taken from a pool of unused, unreserved port addresses) into the source fields of the IP and UDP headers. It will next update the CRC (Cyclical Redundancy Check) and make any other necessary changes to ensure data integrity, then will send the packet to the internet. As part of the process, the gateway will update its internal translation table to cross reference the local machine's IP address with the source port address originally reported by that machine, with the new source port address assigned to the internet-bound packet, and with the destination IP address. Upon receipt of a reply from the internet, the gateway will recognize its own IP address in the packet header, and will examine the incoming packet's destination port address. If it finds the destination port address in its internal table, the gateway will substitute the cross referenced local machine's IP address and original port address into the destination fields of the packet, will update the CRC and any other necessary parameters, and then will

1 dispatch the packet to the LAN, where it will be received by the local machine and directed
2 to the appropriate process. In this manner, a number of computers on a LAN having only
3 local IP addresses can communicate across the internet through a single globally unique
4 IP address.

5 Although NAT gateways provide firewall security against direct accessing of the LAN
6 from the internet, they do not provide security against interception or modification of a
7 packet intended for the LAN while in transit on the internet, and they do not ensure
8 "trustworthiness" from challenges originating within the LAN. Thus, the security provided
9 by IPSec is a necessary protection for LANs that must maintain security while interfacing
10 with the internet.

11 A common implementation of IPSec is to provide security for VPNs consisting of at
12 least one main computing site and one or more remote LANs. The main site and remote
13 LANs are connected across the internet, using that high speed medium to communicate
14 between sites instead of significantly more expensive private leased lines. The drawback
15 to using the internet as a transmission medium, however, is that the internet is inherently
16 insecure, and provides little or no inherent protection against the snooping, detection,
17 "spoofing," or ultimate theft, modification or diversion of messages by hackers. Thus, there
18 is a need for comprehensive security measures where secure data transmissions are
19 required. The IPSec protocol implements security measures to ensure authentication of
20 data and data integrity.

21 The IPSec protocol suite implements security at the network layer of the multi-
22 layered OSI (Open Systems Interconnection) network reference model. The suite includes
23 a number of separate protocols that are used in conjunction with one another to ensure the

1 security of UDP datagrams that carry information across the internet. The base
2 architecture of IPSec compliant systems is explained in RFC2401, "Security Architecture
3 for the Internet Protocol," S. Kent and R. Atkinson (November 1998). The AH
4 (Authentication Header) protocol assures data integrity, source authentication, and
5 incorporates "anti-repeat" measures to deter denial-of-service attacks. ESP (Encapsulation
6 Security Payload) protocol provides protections similar to AH, but adds the additional
7 feature of payload encryption. Both AH and ESP headers have a field for a Security
8 Parameters Index (SPI). The SPI is a 32-bit pseudo-random value that is used to identify
9 a Security Association (SA) for the datagram. Further information regarding these
10 protocols may be found in RFC1826, "IP Authentication Header," by R. Atkinson (August
11 1995), and RFC2406, "IP Encapsulating Security Payload (ESP)," S. Kent and R. Atkinson
12 (November 1998). ISAKMP/Oakley (Internet Security Association and Key Management
13 Protocol, also commonly referred to as Internet Key Exchange – IKE) is a handshaking
14 protocol that establishes the parameters for a secure session between two hosts and
15 provides for the exchange of keying and other security information that is used to
16 implement the secure session and permit the transmission of encrypted data. The
17 ISAKMP/Oakley protocol (hereafter referred to simply as ISAKMP) involves the initial
18 exchanges of unencrypted messages to provide both machines with initialization data from
19 which authentication may be established and secure keys for data encryption may be
20 generated. An explanation of these processes may be found in RFC2409, "The Internet
21 Key Exchange," D. Harkins and D. Carrel (November, 1998). Once security parameters
22 sufficient establish Security Associations (SAs) between hosts have been exchanged, all
23 subsequent transmissions will be encrypted and fully authenticated in accordance with the

1 agreed-upon protocols. At that point the ISAKMP protocol terminates. Subsequent
2 addressing is based upon the IP address for each machine and the machine's SPI for that
3 session. The SPI is unique for each machine during a session. The gateway for a private
4 LAN will maintain an internal table in which "SPI-in" is a value that is cross-referenced to
5 the local machine's IP address, and "SPI-out" is cross-referenced to the IP address of the
6 machine on the internet that is communicating with the local machine. The SPI for each
7 machine is computed from information exchanged during the ISAKMP transmissions, and
8 is carried in the AH or ESP header that is appended to UDP packets. Because IPsec
9 protocols may be nested to provide security in a variety of environments, a single datagram
10 may include both an AH and an ESP header, and may encrypt some header information.

11 Each of the foregoing security protocols modifies the UDP packet by placing new
12 header information on the packet, modifying certain fields within the packet to conform to
13 the protocol being used and, in some cases, encrypting the payload and all or parts of
14 other packet headers. Thus, under IPsec, when a UDP datagram leaves a "secure"
15 domain for transit across an untrusted network, it will normally consist of an IP header, an
16 AH or ESP header (or both), and an encapsulated payload. Header information will include
17 a destination address, an SPI, and sufficient SA information to ensure that the datagram
18 reaches its destination and can be authenticated to the destination host. Encapsulation
19 of the payload ensures that information contained within the payload is denied to unwanted
20 eavesdroppers and hackers. The initial destination host for the datagram may be a router,
21 gateway, or firewall between a LAN and the internet. Upon arrival at the device on the
22 border between the LAN and the internet, the datagram may be opened, examined or

1 decrypted in whole or in part, analyzed for further address information, and routed to a
2 local IP address on the LAN.

3 The ISAKMP handshaking protocol used in IPsec requires that both hosts intending
4 to establish a secure session between them use a process-specific port address (Port 500)
5 for initial message exchanges. For this reason, Port 500 has been assigned for exclusive
6 use with the ISAKMP protocol. By convention, computers attempting to negotiate secure
7 communications parameters by employing the ISAKMP protocol must communicate strictly
8 through each computer's Port 500. That is, ISAKMP messages from either computer must
9 identify Port 500 as both the source and destination port addresses. If either computer
10 receives a packet in which Port 500 is not specified as being both the source and
11 destination, the packet will be discarded.

12 While this protocol provides assurance that two hosts are communicating with each
13 other, it becomes unworkable when one host is located on a LAN that uses local IP
14 addresses and a NAT gateway. For example, Host A, having a local IP address on a
15 remote LAN protected by a NAT gateway, wishes to establish a secure session with Host
16 B, located at a main office computing site. Host A would initiate the protocol by sending
17 an unencrypted UDP datagram to Host B, giving the "destination" as Host B's IP address,
18 and the destination port address as "Port 500." However, when the datagram reaches the
19 NAT gateway connecting the remote LAN to the internet, the gateway will translate the
20 destination port address to an arbitrary port number. Upon the arrival of the datagram at
21 Host B, the ISAKMP protocol will not be recognized, and Host B will not respond. The
22 computers will fail to establish a secure session. Because of this difficulty, it has heretofore
23 been believed that the ISAKMP protocol cannot be used to establish a VPN using a NAT

1 gateway where each computer on the remote LAN uses a local rather than a global IP
2 address.

3 It is therefore an object of this invention to provide a gateway that will permit the use
4 of ISAKMP protocol authentication and key exchanges between a computer having a non-
5 global IP address and a host computer, using the internet as a transmission medium.

6 It is a further object of this invention to provide a gateway that will allow any number
7 of computers on a private LAN using local IP addresses to initiate or receive messages via
8 the internet using ISAKMP protocol.

9 It is another object of this invention to provide a method for employing virtual private
10 networking between two or more LAN sites on the internet, using ISAKMP protocol to
11 initiate secure communications.

12 These and other objects of the invention will become apparent throughout the
13 following description.

14 SUMMARY OF THE INVENTION

15 In accordance with the present invention, a computer using a local IP address on
16 a remote LAN attached to an external network such as the internet through a NAT gateway
17 will use the ISAKMP protocol to exchange keys and establish SAs that will support a
18 secure session under IPSec. For non-ISAKMP traffic, the gateway performs address
19 translation as normal. However, whenever a machine on the LAN originates an ISAKMP
20 protocol message, the gateway will identify the datagram containing port addresses of Port
21 500. Upon encountering such a datagram, the gateway will translate the source IP
22 address, but will not translate the source port address, leaving it at Port 500, and will
23 dispatch the packet to the internet with Port 500 designated as both the source and

1 destination port addresses. The gateway will also update its internal table to "bind" Port
2 500 to the local IP address and associate that binding with the external IP address of the
3 destination machine for a predetermined length of time. If a valid reply is not received
4 within the predetermined length of time, the "binding" between Port 500 and the local IP
5 address will be released. This feature is necessary to ensure that Port 500 is not tied-up
6 indefinitely, such as, for example, in a situation in which an ISAKMP protocol transmission
7 has been initiated to an incorrect destination IP address. Under those conditions, the
8 gateway would never receive a valid reply. If there were no timer to release Port 500 after
9 a period during which a valid reply is not received, the port would remain bound to the local
10 IP address until the gateway was reset. For most conditions, a period of two seconds
11 should be a sufficient length of time to maintain the binding between Port 500 and the local
12 IP address while awaiting a valid reply.

13 During the time that Port 500 is bound to a local IP address, the gateway will
14 continue normal datagram processing of datagrams not having Port 500 port addresses
15 while awaiting a valid reply. A valid reply will be a datagram having a source IP address
16 that is the same as the external IP address that is associated with Port 500, and will have
17 both the source and destination port addresses as Port 500. While awaiting a valid reply,
18 the gateway will ignore other UDP datagrams from the external network having Port 500
19 source and destination port addresses, but not the proper source IP address. Also, while
20 Port 500 is bound to a local IP address, datagrams received from the LAN having source
21 and destination port addresses of Port 500 will undergo "normal" address translation in
22 which the Port 500 source port address will be translated to an arbitrary, unused port
23 address before being sent to the external network. Because such a datagram does not

1 have both a source and destination port address of Port 500, it is not a valid ISAKMP
2 datagram, and will be ignored upon reaching its IP destination. If the period of binding Port
3 500 to a local IP address should expire without a valid datagram having been received by
4 the gateway, the binding will be released, and Port 500 will become available for use by
5 the next datagram having a Port 500 source and destination port address.

6 While Port 500 is bound, upon receiving a valid reply datagram having source and
7 destination port addresses of Port 500 and the correct source IP address, the gateway will
8 process the datagram by substituting the IP address of the local machine into the datagram
9 header's destination IP address field, then will pass the datagram through to the LAN for
10 delivery to the local machine. When the datagram leaves the gateway, the gateway will
11 release the binding between the local IP address and Port 500, and will resume normal
12 datagram processing.

13 If a reply having the proper source IP address and port addresses of Port 500 is not
14 received from the external network, the gateway will time-out after a short predetermined
15 length of time. If the gateway should time-out before a valid reply is received, that ISAKMP
16 message exchange cannot be completed, but must be re-initiated.

17 Once the ISAKMP protocol has been completed, and an encrypted secure session
18 is under way, the gateway will perform local address translations by referencing the SPI
19 in the ESP header of incoming and outgoing datagrams. The gateway will also ensure that
20 each packet type (type 50 for an ESP packet) is correct for the datagram being passed
21 through the gateway. Occasionally, a secure session across a VPN will be interrupted, or
22 a new session started. The gateway's first indication of this will be its receipt of a type 50
23 datagram in which the IP addresses are recognized but the SPI associated with the

1 destination does not appear in its internal table. When this happens, the gateway will
2 dispatch the datagram to the destination IP address using the new SPI, and will also set
3 the destination SPI value (SPI-in or SPI-out, depending upon the direction of the
4 transmission) in its table to the new value, and the source's SPI value to zero. Upon
5 receiving a reply to the transmission, the gateway will replace the zero in the SPI field table
6 with the new SPI for the destination IP address.

7 Because the gateway of this invention does not encrypt or decrypt messages, but
8 simply passes the payload (which may be encrypted or unencrypted) through to the LAN
9 or to the internet for processing at the receiving machine, it does not require intensive
10 processing functionality, and may be used for private LANs in which expense and simplicity
11 of setup and maintenance are considerations.

12 A BRIEF DESCRIPTION OF THE DRAWINGS

13 Further objects and advantages of the present invention can be found in the detailed
14 description of the preferred embodiment when taken in conjunction with the accompanying
15 drawings in which:

16 Figure 1 depicts a virtual private network in which a remote LAN using local IP
17 addresses is networked with a main computing site via an external network which may be
18 the internet. The LAN is connected to the external network through a NAT gateway.

19 Figure 2 depicts a decision chart used by the gateway of this invention to process
20 UDP datagrams received from the LAN for transmission to the internet.

21 Figure 3 depicts a decision chart of steps used by the gateway of this invention to
22 process UDP datagrams received from the internet for delivery to a device on the LAN.

Figure 4 is provided for reference in following the charts shown in Figures 5, 6 and 7. Figure 4 is a table containing IP addresses of local machines on a LAN (L-1 through L-3), the internal and external IP addresses of a gateway, and the IP addresses of external devices ("targets" T-1 through T-3) on an external network.

Figures 5a - 5c show representative fields from an internal table of the gateway cross referencing local IP addresses of machines on a LAN (L-1, L-2, . . . L-x) and external IP addresses of external devices (T-1 through T-3) with SPIs (security parameter indexes) used to authenticate encrypted datagrams. SPI-out represents the SPI of an encrypted datagram leaving the gateway for a device on the internet, while SPI-in represents the SPI of an encrypted datagram destined for a local machine on the LAN. Each view of the table, a, b and c, reflects header values for source, destination, and SPI at different points in time. The changing values signify the commencement of a new session by one a local machine with a target machine.

Figure 6 shows representative fields in datagram headers being exchanged between a single local machine and a single device on the external network. Header values are modified through processing by the gateway of this invention.

Figure 7 shows representative fields in datagram headers being exchanged between three local machines (L-1 through L-3) on a LAN, and three targets (T-1 through T-3) on an external network as they are modified through processing by the gateway of this invention.

Figure 8 is a schematic diagram of signals passing between the datagram processing function and the timer.

DETAILED DESCRIPTION OF THE DRAWINGS

In Figure 1, a virtual private network (VPN) is shown in which a private local area network (LAN) 10 is connected to a computing site 30 located on the internet 50. The LAN 10 uses local IP addresses, and is connected to the internet through the network address translation (NAT) gateway of this invention 20. The computing site 30 may be a business headquarters, or one of any number of private LANs used by a multinational corporation, an educational facility, or any other site which will be frequently accessed from remote locations. Such sites will normally have a firewall or gateway 35 capable of running encryption and other security applications. Such a gateway will have the ability to open a packet, decrypt or otherwise access its contents, and perform address translation, routing, de-encapsulation, and data manipulation functions as well. While such devices are able to support ISAKMP and other IPsec protocols, they do so by opening and decrypting packets, and manipulating data, and are, by and large, too expensive and powerful to be employed efficiently at remote LAN sites needing to establish a VPN with the main computing site.

A server 40 at the main site runs the VPN server software. Computers 15 at remote sites each run appropriate VPN client software that implements IPsec security protocols on each computer.

A computer 15 on the LAN 10 will communicate with devices on or across the internet through gateway 20 by sending an IP datagram to a server 40 at the computing site 30.

Datagrams received at the gateway 20 are processed according to the decision charts shown in Figures 2 and 3. Although the flow charts of Figures 2 and 3 show both

the processing steps and a sequence for the steps, the order for performing some of the functions is not critical, and some of the steps may be done in an order other than is shown in the flow charts without affecting the ultimate result. For example, Figures 2 and 3 show that the first step after a datagram is received by the gateway is to determine the datagram type, while the last step is to perform the IP address translation that is necessary before the datagram is passed through the gateway. Some embodiments, however, could place the step of address translation to some point earlier in the process, and this would not affect the outcome of the process. Since the order of translating the IP address is not critical to the overall process, the determination of when this translation should be done is a matter of engineering choice.

As shown in Figure 2, upon receiving a datagram from the LAN, the gateway will check to see whether the datagram is encrypted. It will do so by checking the "Next Header" field in the IP header to determine the type of datagram it is dealing with, and to see whether the datagram has been encrypted. A datagram type of 50 (ESP) indicates that the datagram is encrypted, and port address information may not be available.

Continuing through the decision tree of Figure 2, if the datagram is encrypted, the gateway will check the datagram's SPI to see whether it appears in the SPI-out field of the gateway's internal table. Representative fields from such a table are shown in Figures 5a - 5c. If the SPI of the datagram is found in the SPI-out field of the internal table, the gateway will modify the source IP address of the datagram to be the external IP address of the gateway, and will send the datagram to the external network for delivery to an external device.

1 If the datagram is encrypted, but the SPI does not appear in the gateway's internal
2 table, then according to the decision chart of Figure 2 the gateway will assume that the
3 datagram is initiating a new session. In this case the gateway will set the SPI-in field of its
4 internal table to zero (0) and will set SPI-out to the new SPI from the datagram. These
5 modifications to the internal table are reflected in Figures 5a and 5b, in which a "new" SPI
6 ("14662"), which did not appear in the SPI-out field of the gateway's internal table in Figure
7 5a, is shown as having been entered into the SPI-out field, and SPI-in has been set to zero
8 (0) in Figure 5b. The encrypted datagram will then be sent to the external gateway after
9 the source IP address has been translated from that of the local device to the external IP
10 address of the gateway. These steps are shown in Figures 5b and 5c.

11 Continuing with the decision chart of Figure 2, if the datagram is not encrypted, the
12 gateway will next check the datagram's destination port address. If the port address is
13 anything but Port 500, the gateway will enter the source port address into its internal table,
14 cross reference it with the (local) source IP address, and will then substitute an arbitrary,
15 unused port address into the source port address field of the IP header. It will also enter
16 the new port address in its internal table, again cross referenced to the (local) source IP
17 address. This process, which is used for unencrypted datagrams not having Port 500 as
18 a port address, shall be referred to as "normal address translation" for datagrams
19 originating on the LAN. These translations are shown in Figure 6, rows 1 and 2. The
20 datagram will then be dispatched to the internet for routing to the destination IP address.

21 In Figure 2, where an incoming datagram's source and destination port addresses
22 are Port 500, the gateway must next check its tables to see whether Port 500 is already
23 bound to an IP address. If Port 500 is free, then the gateway will "bind" Port 500 to the

(local) source IP address of the datagram, will create an association between the port and the (external) destination IP address, and will send a signal to start the internal timer. The gateway will also process the datagram by substituting the gateway's external IP address for the local IP address in the source IP address field. It will not, however, translate the source port address. By suspending the "normal" translation of the source port address, the gateway ensures that the target machine will be able to recognize the datagram as being an ISAKMP datagram. These steps are also shown in Figure 6, rows 5 and 6.

In Figure 2, if an incoming datagram from the LAN has a source and destination port address of Port 500, but Port 500 is already bound to some other local IP address, then the gateway cannot bind Port 500 for the message then being processed. In that event, the gateway will process the datagram "normally," as if it were not an ISAKMP datagram. That is, it will translate the datagram's source port address to an arbitrary number and will translate the source IP address to be that of the gateway's external IP address. The gateway will then send the datagram to the internet, where it will be rejected by the target because it does not conform to an ISAKMP datagram. This event is depicted in Figure 7 at rows 15 and 16.

In figure 3, a decision chart is shown which outlines the steps the gateway will follow in processing datagrams received from the internet. Upon receiving a datagram, the gateway will first check its type and, if the datagram is encrypted, will check to see whether the SPI appears in its internal table. If the SPI is recognized, its destination IP address will be translated to be the IP address of the local device, and the datagram will be passed to the LAN for delivery to the local device. If the SPI is not recognized, the gateway will next check to see whether its SPI-in field corresponding to the datagram's source IP address

1 is zero (0). If SPI-in is zero, the gateway will assume that the datagram is the first reply of
2 a new session and will replace the zero in the SPI-in field with the SPI of the datagram.
3 The gateway will then translate the destination IP address to be the local IP address of the
4 device on the LAN, and will send the datagram to the LAN for delivery. This event is
5 shown in Figures 5b and 5c. In Figure 5b, the SPI-in for local machine L-1 has been set
6 to zero. Upon the gateway's receipt of a datagram from the internet having an SPI of
7 3288, the gateway will not find that SPI in its SPI-in field. The gateway will next check to
8 see whether the SPI-in field is holding a value of zero. Upon determining that SPI-in for
9 local machine L-1 is zero, the gateway will replace the zero with the SPI of the datagram
10 ("3288") and will send the datagram to the LAN. This is shown in Figure 5c.

11 In Figure 3, if the datagram from the internet is not encrypted, the gateway will
12 check to see whether it has a port address of 500. If it does not, the datagram will undergo
13 "normal" address translation for datagrams from the external network, meaning that the
14 local port address and local IP address of the device on the LAN will be substituted into the
15 destination fields of the datagram, and the datagram will be sent to the LAN for delivery.
16 This "normal" address translation for datagrams from the internet is shown in Figure 6 at
17 rows 3 and 4.

18 Again referring to Figure 3, If the datagram does have a port address of 500, the
19 gateway must next check to see whether Port 500 is bound to a local IP address and is
20 associated with the datagram's (external) source IP address. If it is, then the datagram is
21 valid, and will be passed to the LAN after the destination IP address has been translated
22 from that of the external gateway to the IP address of the local device. Upon passing the

1 datagram to the LAN, the gateway will release Port 500. This event is illustrated in Figure
2 6 at rows 7 and 8.

3 If, in Figure 3, Port 500 is bound to a local IP address, and is associated with an
4 external IP address other than that found in the source IP address of the datagram, then
5 the datagram is not valid and will not be further processed by the gateway. This event may
6 be seen in Figure 7 at rows 25 - 31. At rows 25 and 26, local machine L-1 sends an
7 ISAKMP datagram to target T-1. At this point, Port 500 is bound to the IP address of local
8 machine L-1 and is associated with the IP address of target T-1. However, as shown in
9 Figure 7, the timer "times" out before a reply is received at the gateway from T-1, and, at
10 row 27, Port 500 is released. At rows 28 and 29, local machine L-3 sends an ISAKMP
11 datagram to target T-3, binding Port 500 to L-3's IP address and creating an association
12 with T-3's IP address. While Port 500 is bound, a reply is received from T-1. However,
13 because Port 500 is bound, and is associated with T-3's IP address, the reply from T-1 is
14 discarded. This is shown at rows 30 and 31 of Figure 7.

15 Figures 5a - 5c depict an internal table of the gateway in which the IP addresses
16 and SPI numbers for encrypted communications between local computers and targets on
17 the internet are maintained. Fields for "L-1," "L-2," "L-x," and "T-1" through "T-3" are
18 included for ease of reference, and do not appear in the gateway's internal tables. In
19 Figure 5, the field "SPI-out" holds the SPI for each target machine during a secure session
20 with a specific computer on the LAN. The "SPI-in" field gives the corresponding SPI that
21 will be recognized by the local computer as signifying a valid datagram intended for it.
22 Figure 5a shows the table at a beginning time. Eight (8) local computers have participated
23 in encrypted sessions with three targets, T-1 through T-3 during the life of the table's data.

1 This is shown by the fact that each local machine shows an SPI-in associated with its IP
2 address. Although only three targets are shown in the table, it may be noted that each
3 target is using a different SPI-out for communications with each local machine. In this
4 manner, a target will know from which source an encrypted datagram was generated.

5 Figure 5b shows the same local and target computers as Figure 5a. Here, however,
6 the SPI-out for the session between L-1 and T-1 is a new SPI, indicating a new session
7 between the computers. The gateway's first indication that a new session is taking place
8 is its receipt of an encrypted datagram from the LAN having an SPI – "14662" – that is not
9 in its table. The gateway forwards the datagram to the internet, but also modifies its table
10 to place the new SPI in the SPI-out field associated with the source and destination IP
11 addresses for that datagram. It also places a zero in the SPI-in field as a marker to
12 indicate that a new SPI-in is also expected. Figure 5c shows that a new SPI – "3288" –
13 was included in a datagram received from T-1. That SPI has been entered into the
14 gateway's SPI-in field, and further communications between L-1 and T-1 during this
15 session will use those SPI's to authenticate their messages.

16 Figure 6 charts the flow of representative datagrams through the gateway of this
17 invention by a single computer on a LAN communicating with a remote target on the
18 internet. Each row of the chart represents information in a datagram at either the LAN
19 interface with the gateway or the internet interface with the gateway. Consecutive rows
20 represent data entering the gateway from one side and leaving the gateway at the other.
21 The gateway has one IP address, which may be a local IP address, at its interface with the
22 LAN, and a global IP address at its interface with the internet. The columns in Figure 6
23 depict the side of the gateway the datagram is traversing, the type of datagram, the

1 datagram's source IP address and port address, the datagram's destination IP address
2 and port address, and the datagram's Security Parameter Index (SPI) for encrypted
3 datagrams of type 50, using ESP (Encapsulated Security Payload) protocol.

4 Row 1 of Figure 6 shows a UDP datagram arriving at the local interface of the
5 gateway, and having a source IP address corresponding to local computer L-1, and a
6 destination IP address of the target on the internet, T-1. For ease of reading, Figure 4
7 provides a table of IP addresses cross referenced with local designations L-1 through L-3,
8 and target designations T-1 through T-3. The source port address for L-1 is Port 6404, and
9 the target's destination port is Port 80. Since the datagram is not encrypted, and does not
10 exhibit a port number of 500, it undergoes a normal translation in which an "arbitrary" port
11 address, Port 10425 is substituted into the source port address field and the gateway's
12 external IP address is substituted for the source IP address of the datagram. Although the
13 translated source port address is said to be "arbitrary," it will normally be the next in a
14 sequence taken from a pool of unreserved and presently unused port addresses
15 maintained by the gateway.

16 As the datagram exits the gateway, as shown Figure 6, in row 2, the address
17 translation function of the gateway has substituted the gateway's external IP address into
18 the datagram header for the source IP address, and has given the source port an arbitrary
19 number. Rows 3 and 4 show the reply datagram from the target. In row 3, a UDP
20 datagram from the target shows the destination IP address as being the external IP
21 address of the gateway, and a destination port as being the port address arbitrarily
22 assigned by the gateway. Since the datagram is not encrypted and does not have a port
23 address of 500, the datagram undergoes normal translation of the destination port address

1 and IP address, then is sent to the LAN. In row 4, the gateway has substituted the local
2 computer's local IP address and port address in the destination fields of the header before
3 sending the datagram to the LAN.

4 In row 5 of Figure 6, the local computer initiates an ISAKMP protocol with the target.
5 The datagram type is shown as ISAKMP. Both the source and destination port addresses
6 are Port 500. When the gateway determines that the destination port address is Port 500,
7 it checks to see whether Port 500 is currently bound to any IP address. Since it is not, the
8 gateway passes the datagram, translating only the source IP address field to show the
9 gateway's external IP address, but without changing the source port address.

10 In Figure 6, rows 5 - 16 show the six standard ISAKMP "handshaking" datagram
11 exchanges necessary to establish SAs (Security Associations) to support fully encrypted
12 and authenticated datagrams. Although some modes of ISAKMP use fewer exchanges,
13 the main mode is depicted in Figure 6. Following the establishing of SAs, the local
14 computer and the target begin communicating using ESP protocol encrypted datagrams.
15 Here, datagram validity is maintained through the use of Security Parameter Indexing –
16 SPI – numbers in an SPI field of the datagram's header. Each host recognizes a datagram
17 "addressed" to its SPI, which can be modified during a session by mutual agreement of the
18 hosts as necessary to ensure continued security. When an encrypted datagram passes
19 through the gateway, as depicted at Figure 6, rows 17 and 18, neither the source nor the
20 destination SPI is modified by the gateway, although the datagram's source IP address is
21 translated to be the gateway's external IP address.

22 Thus, when an encrypted datagram is received by the gateway, it will be signified
23 by a datagram of type 50 (ESP). Upon encountering that datagram type, the gateway will

1 check the datagram's Security Parameter Index (SPI) to see whether that SPI is recorded
2 in its internal table. If it is, the gateway will translate the datagram's source or destination
3 IP address, as appropriate, and will send the datagram to the LAN or the internet,
4 depending upon the direction of transmission. However, if the SPI of a datagram from the
5 LAN does not appear in the gateway's internal table, and the source and destination are
6 recognized IP addresses, the gateway will assume that a new session has been started.
7 In this case, it will pass the datagram to the external network leaving the new SPI intact,
8 but recording the new SPI in the "SPI-out" field of its internal table and placing a zero into
9 the "SPI-in" field. At rows 25 and 26 it may be seen that a new SPI has appeared,
10 signifying a new session. This event corresponds to figure 5b, where the "0" in the "SPI-in"
11 field corresponds to the new SPI-out of "14662." At rows 27 and 28, the reply packet from
12 the external network shows that "old" SPI "9802" has been replaced with "new" SPI "3288."

13 Figure 7 is similar to Figure 6, except that it illustrates the passage through the
14 gateway of this invention of datagrams between three computers on a LAN, designated L-
15 1, L-2, and L-3, and three targets on the internet having unique global IP addresses, T-1,
16 T-2 and T-3. In Figure 4, for ease of reference, a table containing the IP addresses of
17 these devices is given. As shown in Figure 7, a transmission designated "L-1 Out"
18 represents a transmission from local computer L-1 to the gateway. "T-1 In" represents a
19 transmission from the gateway to target T-1. "T-1 Out" represents a transmission from
20 target T-1 to the gateway, while "L-1 In" represents a transmission from the gateway to
21 computer L-1.

22 As shown in rows 1-8 of Figure 7, computers L-1 and L-2 conduct "in the clear"
23 communications with targets T-1 and T-2. At row 9, L-1 commences an ISAKMP session

1 with T-1. Rows 9-14 show the first three messages exchanged between L-1 and T-1
2 during the ISAKMP protocol. At row 15, computer L-3 commences an ISAKMP-1 message
3 exchange with T-3. However, at that time Port 500 is bound to L-1 and is associated with
4 the IP address of T-1, awaiting an ISAKMP-4 reply from T-1. In this situation, the datagram
5 from L-3 cannot bind Port 500, and its source port address will be translated. As such, L-3
6 cannot complete the transmission that was started at row 15.

7 Thereafter, at rows 17-18, T-1's reply (ISAKMP-4) is received at the gateway and
8 sent to L-1, and Port 500 immediately becomes available. Thus, when L-3 reattempts its
9 ISAKMP-1 transmission at row 19, the transmission is successful.

10 At rows 19-20 of Figure 7, L-3's ISAKMP-1 transmission binds Port 500 to L-3's IP
11 address. Thus, when L-1 attempts its ISAKMP-5 transmission, at rows 21-22, Port 500 is
12 not available, and the gateway simply translates the destination port address from Port 500
13 to an "arbitrary" port number – in this case, "9063" – and sends the datagram to the
14 internet, where target T-1 will not recognize it as an ISAKMP datagram. However, after L-3
15 releases Port 500, at rows 23-24, L-1's next attempt to send its ISAKMP-5 transmission
16 is successfully received by T-1. However, T-1's reply is slow, and, at row 27, Port 500 is
17 released from its binding to L-1, and, at rows 28-29, is immediately grabbed by L-3 for an
18 ISAKMP-3 transmission. Thus, when T-1's ISAKMP-6 reply arrives at the gateway, as is
19 shown at rows 30 and 31, Port 500 is blocked, and the datagram is ignored. Thereafter,
20 L-1, not having received a reply to its ISAKMP-5 message, retransmits it at rows 34-35,
21 and a reply from T-1 is received at rows 36-37. Following their ISAKMP handshaking, L-1
22 and T-1 can communicate securely, using ESP protocol at rows 38-39 and 42-43.

1 Rows 38-57 of Figure 7 demonstrate the gateway's handling of a variety of
2 datagrams between a number of local computers and targets. UDP datagrams are shown
3 at rows 40-41, ESP datagrams at rows 42-43 and 52-53, and ISAKMP datagrams at rows
4 44-45. While the chart of Figure 7 shows different IP addresses for each device, in
5 practice it may occur that a number of processes will be running on the same device. The
6 substitution of unique source ports by the gateway, and the use of SPI's to differentiate
7 encrypted transmissions ensures that datagrams emanating from multiple processes
8 running on a single machine will not be misdirected.

9 Figure 8 depicts the initiation and transfer of signals between the datagram
10 processing circuitry 100 and the timer 110. Upon the occurrence of an event requiring a
11 port address to be bound to an IP address, a signal 120 will be sent to the timer to
12 commence timing. Upon the expiration of the appropriate interval, the timer will send a
13 signal 140 indicating that time has expired, in which case any port that is bound will be
14 released. In the interim, if an expected datagram has arrived, and a previously bound port
15 is to be released, a disabling signal 130 will be sent to the timer indicating that the timer
16 should be reset to await the next signal to begin timing. Obviously, there are numerous
17 timing circuits known in the art, and the specific configuration shown in Figure 8 is only one
18 of many possible embodiments.

19 From the foregoing it will be understood by those of skill in the art that the preferred
20 embodiment described herein is not the only means for practicing the invention, and that
21 other embodiments may be chosen to practice the invention without departing from the
22 spirit and scope of the invention. For example, although the preferred embodiment is
23 described with reference to Port 500, which has been reserved exclusively for use with the

1 ISAKMP protocol, the invention may be employed in the same manner for processing
2 datagrams destined for other port addresses that may in the future be assigned to other
3 processes or protocols. In particular, many games played over the internet require the use
4 of specific ports on local and external machines that cannot withstand normal address
5 translation. Additionally, although the invention has been described primarily with respect
6 to communications between a private LAN and the internet, it will be apparent that the
7 gateway of this invention can be used at any interface between two networks and will have
8 the same functionality as has been described.

9 The claims appended hereto are meant to cover modifications and changes within
10 the spirit and scope of the present invention.

11 What is claimed is: